

Viewpoint Paper ■

Why Do It the Hard Way? The Case for an Expressive Description Logic for SNOMED

ALAN L. RECTOR, MD, PhD, SEBASTIAN BRANDT, PhD

Abstract There has been major progress both in description logics and ontology design since SNOMED was originally developed. The emergence of the standard Web Ontology language in its latest revision, OWL 1.1 is leading to a rapid proliferation of tools. Combined with the increase in computing power in the past two decades, these developments mean that many of the restrictions that limited SNOMED's original formulation no longer need apply. We argue that many of the difficulties identified in SNOMED could be more easily dealt with using a more expressive language than that in which SNOMED was originally, and still is, formulated. The use of a more expressive language would bring major benefits including a uniform structure for context and negation. The result would be easier to use and would simplify developing software and formulating queries.

■ *J Am Med Inform Assoc.* 2008;15:744–751. DOI 10.1197/jamia.M2797.

Introduction

Since its first release in 2002, several countries around the world have embraced SNOMED-CT (SNOMED for short) as a reference terminology for their national health care institutions. Apart from changes and extensions to content, however, neither the structure of SNOMED nor the expressiveness of the underlying formalism—OntoLog—has changed significantly since their initial development in the early to mid 1990s.

Since then, there have been significant developments in both logic-based formalisms and ontology design. Researchers have reviewed SNOMED in the light of these advances, and several proposals for improvements have been made. For example, Bodenreider¹ examined the specialization hierarchy of SNOMED classes and suggested that thousands of classes are apparently defined at variance with basic ontological principles. Schulz² discussed ‘relationship groups,’ a construct unique to SNOMED's representation language. He suggests that relationship groups can be replaced by mereological relations using syntactic constructs available in all modern ontology languages, including OntoLog, and that this would significantly improve the clarity of the definitions. In a separate publication by Schulz and his colleagues,³ a broad range of ontological problems in SNOMED are identified and a comprehensive set of remedies are

proposed, affecting not only the definitions of classes but also the relations used in SNOMED. In that paper, the authors advocate a modest extension of the logical formalism underlying SNOMED, arguing that it would permit simpler and clearer definitions of classes and, especially, of the relations that modify and link classes.

In this paper we review these problems in general but focus in particular on three issues: SNOMED's “context model” and the notion of “Situations involving specific context,” the representation of part-whole relations, and the problems of determining semantic equivalence between findings and observables. Following on from the theoretical arguments in previous papers,^{4,5} we argue for a schema that integrates context with other concepts, so that all related concepts appear in the same hierarchy rather than (as currently) potentially in two parallel hierarchies, one for the concepts themselves, and one for the concepts occurring in situations. Such an integrated schema would require extending SNOMED's logical formalism further than proposed by Schulz, to one that includes negation, disjunction, and “general concept inclusion axioms (GCIs)”;⁶ the obvious candidate for such a formalism is the W3C standard Web Ontology Language (OWL 1.1). Although in the past the scale of SNOMED has been a barrier to the use of OWL and related formalisms, this is no longer the case for the existing schema, and it seems likely that a reformulation using more expressive schemas would likewise prove tractable. At the same time, these proposals seem likely to improve SNOMED's “cognitive scaling”—the cognitive load facing authors dealing with a very large corpus. Given the potential advantages, it is important to test the hypothesis that the more expressive schemas proposed here will prove scalable.

We argue that reformulation using such a schema and a more expressive language would have major advantages:

- A uniform, clear, and understandable schema for all concepts used in clinical records including negation and context.

Affiliation of the authors: University of Manchester, School of Computer Science, Manchester, UK.

This work was supported in part by the UK Department of Health “Connecting for Health” programme, the UK MRC CLEF project (G0100852), the JISC and UK EPSRC projects CO-ODE and HyOntUse (GR/S44686/1), the EU funded Semantic Mining Network of Excellence and SemanticHEALTH FP6 Specific Support Action, IST-27328-SSA, <http://www.semantichealth.org>. The HL7 Terminology working group stimulated and contributed to many of the ideas presented here.

Correspondence: Alan Rector, School of Computer Science, University of Manchester, Oxford Road, Manchester M13 9PL, England; e-mail: <rector@cs.manchester.ac.uk>.

Received for review: 03/17/08; accepted for publication: 07/25/08.

- Elimination of the need for special mechanisms to deal with context, partonomy, and role groups.
- More effective leveraging of the underlying logical representation to organize and quality assure the SNOMED hierarchies.
- Improved ability to recognize semantic equivalence between post-coordinated and pre-coordinated expressions and between “observables” with “values” and the corresponding “findings.”
- Improved ability to modularize and segment SNOMED for specific purposes.
- Access to the tools and techniques being developed by the wider Semantic Web and OWL communities.

The overall result is that SNOMED would be more regular, uniform, and have a better defined and more consistent semantics. This in turn would make it easier to use, query, quality assure, and use as the basis for software.

In outline, the proposals are:

- To represent all concepts used in clinical records (findings, observables, and procedures) uniformly as fully defined “situations” that include any context required and that deal with negation explicitly and formally.
- To represent all anatomical sites and lesions so that it is explicit whether the structure is intended to include only the entity in its entirety or whether it is to include the parts of the entity as well—e.g., whether the class of procedures “removal of lung” is to include only removal of an entire lung or if it is intended to include as well removals of lobes and segments of lungs.
- To define observables and related findings in such a way that the classifier can be used to recognize the equivalence between an observable with a given value and the corresponding finding of the observable with that value—e.g., between an observable of “blood pressure” qualified by “increased” and a finding of “increased blood pressure.”
- To organize SNOMED as a set of modules that can be easily separated for specific applications.

In the following section, we begin by highlighting features of SNOMED that might benefit from reformulation using established OWL modelling patterns. Then, in the discussion, we discuss the feasibility of adopting the proposed changes to SNOMED, their potential benefits, and unresolved issues.

SNOMED from an OWL Perspective

The OWL 1.1^a is strictly more expressive than Ontology, and tools exist to load Ontology directly into OWL 1.1^b. For uniformity, therefore, we present both existing SNOMED definitions and proposed extensions in OWL 1.1 using the Manchester syntax for readability.⁶ Moreover, since SNOMED “Concepts” are equivalent to OWL “Classes” we use the two terms interchangeably.

Situations with Explicit Context

Within SNOMED, findings and conditions can either appear as plain subconcepts of clinical finding, condition, etc., or can be embedded within a construct previously called

```
391094005 H/O: vertebral fracture equivTo:
(243796009|Situation with explicit context|:
  (246090004|Associated finding|= (64572001|Disease|:
    (116676008|Associated morphology|= 72704001|Fracture|,
    363698007|Finding site|= 51282000|Bone structure of spine|)),
  408729009|Finding context|= 410515003|Known present|,
  408732007|Subject relationship context|= 410604004|Subject of record|,
  408731000|Temporal context|= 410513005|Past|))
```

Figure 1. Definition of history of vertebral fracture in SNOMED.

“context dependent concept”, now renamed “situation with explicit context.” “Situations with explicit context” can be used to specify additional information relating to:

- The presence or absence of the phenomena under consideration;
- “Modalities” such as “risk”, etc.;
- “Temporal” positioning such as “past,” “present actual,” etc.;
- “Subject of care,” such as “subject of record,” “fetus,” etc.

For instance, history of vertebral fracture is defined in SNOMED as shown in Figure 1.

The class does not merely define vertebral fracture, but a fracture known to have been present in the bone structure of the spine of the subject of the record in the past. The way these situations are defined in SNOMED has three major drawbacks.

Firstly, a finding of interest can either occur on its own, e.g.,
195826005 | Nasal obstruction,
or as a situation:

267100006 | Nasal obstruction present (situation).

This means that any query for the notion of nasal obstruction must look in both places and that the distinctions between the two hierarchies are often unclear.

Secondly, because the Ontology formalism does not support negation, negation is expressed by qualifiers for “presence” and “absence” that require special procedures for querying and classification. This makes it difficult to be certain of the intended meaning or to verify that situations involving negation are classified correctly.

Thirdly, some of the additional qualifiers available for the definition of situations conflate categories that would be better dealt with separately—e.g., “presence/absence” and “history of.” The result is that there are many special cases and the model has become over-complex, as indicated, for example, by the scale of the detailed documentation required for the use of SNOMED in HL7.⁷

Moreover, a careful review suggests that no situation with explicit context defined in SNOMED includes the absence or presence of more than one proper clinical finding. Apparent exceptions are cases where the associated finding also carries temporal information or is altogether redundant. For instance, 407553003 | History of glandular fever has the associated findings 40733004 | Infectious disease (disorder) and 307294006 | Personal history finding, the latter rather addressing the temporal context. In 161496006 | History of chronic ear infection, the findings associated are 129127001 | Infection of ear and 118236001 | Ear and auditory finding, the latter being redundant due to the

^a<http://www.webont.org/owl/1.1/>.

^bProtege4 available from <http://protege.stanford.edu>.

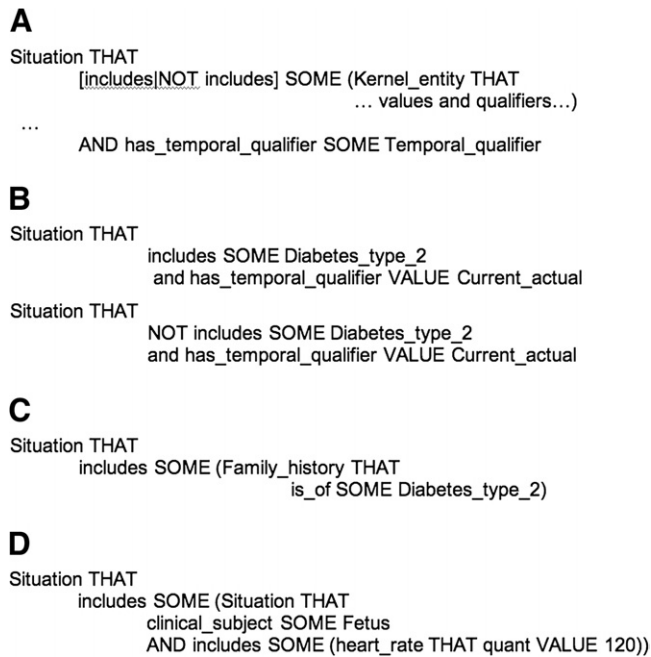


Figure 2. A. Proposed skeleton of the schema for a Situation in OWL. B. Examples of situations representing the presence and absence of “Type 2 Diabetes”. C. Example of situation representing “Family history of type 2 diabetes”. D. Example of situations representing “Fetal heart rate of 120”.

definition of the former. Without clear formal semantics, clarifying such issues is, at best, difficult.

Given the flexibility of class definitions in OWL, the above drawbacks are straightforward to overcome. Firstly, the distinction between findings within and without situations can be removed by separating concepts into “kernel concepts” that represent the entity itself and “recordable concepts” for situations in which the entity either is, or is not, present. Kernel concepts would then be purely internal. Only recordable concepts would correspond to codes used as statements in Electronic Health Records (EHRs).

Each recordable code would represent a class expression for a “Clinical situation,” at its simplest, the skeleton schema is shown in Figure 2a and examples shown in Figure 2b. Presence or absence is dealt with by the choice of either “includes” or “NOT includes”. Unlike the current formulation in Ontylog, “NOT includes” is a formal logical construct and is dealt with automatically by the classifier without further special treatment.

Qualifiers such as “risk” or “family history” that, in SNOMED parlance, “modify the axis” would be in the modelled situation differently from other qualifiers, as “prefixes,” e.g., as in Figure 2c.

Family history, risks, etc. can be either included or not included—i.e., be stated explicitly to be either present or absent just as for kernel concepts. Because of their placement as “prefixes,” all danger of confusion of axes is eliminated.

Using a related mechanism, subjects other than the patient can be dealt with by nesting situations, e.g., by nesting a situation about the fetus within a situation about the mother as shown in Figure 2d.

Classes and Their Definitions

Although ontology languages such as Ontylog and OWL can be used merely to construct hierarchies of classes manually, the true strength and purpose of such logic-based languages is that they allow classes to be defined by complex class expressions built recursively from previously defined classes and properties using constructors provided by the ontology language. The benefits of this approach are that (1) the meanings behind the classes are made explicit, (2) the actual hierarchy of classes can be computed automatically on the basis of their definitions in the ontology, and (3) multiple hierarchies can coexist or be extracted for different use cases.

In the 2008 SNOMED release, only 16% of the classes are fully defined. Furthermore, the vast majority of the remaining “primitive” classes have only ‘trivial’ information asserted about them.¹ The information is usually insufficient even to permit consistency checking, let alone automatic classification.

One reason for the limited use of fully defined classes in SNOMED is that, in Ontylog, once a class is fully defined, it cannot be further qualified. This limits the use of fully defined classes to those which need not be further qualified.

More technically, in Ontylog, *fully defined* classes are introduced by expressions using the Keyword “*defconcept*” and specify set of necessary and sufficient conditions. Expressions using “*defconcept*” correspond to OWL “*equivalentClasses*” axioms (demoted in the Manchester OWL syntax by “*EquivalentTo*” and abbreviated in this paper to “ \leftrightarrow ”). *Primitive classes* are introduced by expressions using the keyword “*defprimclass*” and introduce a set of necessary but not sufficient conditions. Expressions introduced by “*defprimconcept*” correspond to OWL “*subclassOf*” axioms (denoted in the Manchester OWL syntax by “*subclassOf*” and abbreviated in this paper to “ \rightarrow ”). In Ontylog, the same class cannot be the subject (left-hand side) of both a *defconcept* statement (*equivalentClasses* axiom) and a *defprimconcept* statement (*subclassOf* axiom). In OWL, there is no such limitation. The result is that, in OWL, additional information can be added to fully defined classes. In fact, this is just a special case of OWL’s support for what are termed “general inclusion axioms”—i.e., axioms that allow any class, primitive or defined, to be asserted to be a subclass of any other class, including the “restrictions” that correspond to SNOMED qualifiers.^c

Since in our proposed schemas, all “recordable concepts” correspond to fully defined classes, OWL’s support for “general inclusion axioms” is essential. Given general inclusion axioms, the proposed schema leads straightforwardly to a unified classification hierarchy in which all codes that include the presence of a given concept occur together. The need to search in two parallel hierarchies for every concept is eliminated and, furthermore, alternative classifications—e.g., of all terms involving the family history, past history, or current diabetes—can be produced simply using automatic classification.

^cTechnically, qualifiers in SNOMED correspond to “restrictions” in OWL. A “restriction” in OWL is just a special kind of class, the class of all those entities that satisfy the restriction.

Representing Parts and Wholes

Any clinical representation needs to support the pattern that, in general but not always, a disorder of the part is a disorder of the whole. For example, a disorder of a heart valve is a disorder of the heart; a fracture of the neck of the femur is a fracture of the femur; a procedure on a lobe of the lung is a procedure on the lung; etc. The SNOMED originally dealt with this issue using “right identities,” a special type of axiom for properties. For example, the right-identity statement “site *o* is_part_of \rightarrow site” expresses the rule that it is always the case that anything that has a site that is a part of a whole also has the whole as a site. Right-identities are equivalent to the construct that GALEN called “refinement”⁸ and can be expressed in OWL 1.1 by a more general construct called “property chains.” (There is an extensive literature on parts, wholes, and sites. For discussions of the issues from different points of view see the reference citations 9–12.)^{9–12}

Using property chains (or SNOMED’s more limited “right identities”) to represent this pattern, however, requires great care. For instance, the valves are part of the heart, yet failure of a heart valve does not imply heart failure, although it may lead to it. Similarly, although we may want the notion of “Lung operation” to include “removal of a lobe of the lung,” we would not want “pneumectomy” (removal of the lung) to include “removal of a lobe of the lung.” These problems originate because our language often relies on “common sense” to make these distinctions, so that a literal representation of the language may not give the desired result.

Recent versions of SNOMED have experimented with an approach suggested in Schulz and Hahn, 2001¹³ involving so-called SEP-Triples—triples of the thing in its entirety (E), its parts (P), and the disjunction of the thing and its parts (S). This approach was found cumbersome when implemented literally because it necessitated enumerating three nodes for every anatomical structure.^d

As OWL supports disjunctions, the distinction between whether a property applies to just the thing itself, just its parts, or both is easy to express. See Figure 3 for an example.

The first class expression includes all operations on a lung as a whole or any of its parts; the second expression includes the removal of a lung as a whole; the third expression represents the removal of a lobe of a lung which, since it is a part of a lung, will be classified under the first expression automatically. The same mechanism deals naturally with “partial” and “total,” e.g., “Removal THAT site SOME Kidney” represents a total nephrectomy; “Removal THAT site SOME (is_part_of SOME Kidney)” represents a partial nephrectomy. The disjunction—Removal THAT site SOME (Kidney OR is_part_of SOME Kidney)—includes both total and partial nephrectomies. This achieves the same distinctions as possible with SEP triples but without having to create the additional nodes explicitly.

Observables and Findings

One of the persistent issues in clinical information systems is the distinction between observables and findings. Although there exists no universal consensus on the distinction, the term “observable” generally refers to an aspect of the patient that can be quantified or qualified, e.g., “blood pressure,” “skin color,” “body-mass index,” etc. A “finding,” on the other hand, usually refers to something which is either present or absent, possibly with additional qualification, e.g., “diabetes,” “fractures,” etc., or to the state of some observable such as “increased blood pressure” which likewise may be present or absent.

In SNOMED, distinctions are made between the classes “finding” and “observable entity” but the relationship between them is not always easy to understand. As an example, consider the finding of increased blood pressure defined in SNOMED as shown in Figure 4a. Hence, the finding of increased blood pressure implies a finding of “abnormal blood pressure” that interprets the observable entity “blood pressure.” The fact that a finding of an “increased blood pressure” qualifies the blood pressure as abnormally *high* as opposed to abnormally low is not reflected at all in the expression. This is a common phenomenon. In many cases, most of the intended meaning behind concepts such as finding of increased blood pressure remains in the term name and is not reflected in a definition. This is even more obvious when comparing SNOMED’s (primitive) definition of a decreased blood pressure as shown in Figure 4b.

A comparison shows that there is no distinction between the definition of increased and decreased blood pressure with respect to the actual blood pressure value observed. Instead the definitions are different because a Finding of increased blood pressure implies an abnormal blood pressure whereas a Finding of decreased blood pressure does not. (Whether or not this is intentional cannot be determined from the information available.)

Furthermore, in many models of the medical record, it is possible to express “increased blood pressure” by the combination of the code for “blood pressure” and a code for “increased” or some equivalent, which in SNOMED might be a qualifier such as 75540009|high| or 260399008|raised|. However, because no such qualifier is present in the SNOMED definition of “increased blood pressure,” it is not possible to recognize the equivalence between the two formulations.

Using GCI’s available in OWL ontologies, it is easy to represent the relation between observables and findings more faithfully. If the definition of the kernel concept for increased blood pressure finding were to be as shown in Figure 5a.

Then whether or not the named finding were used, the underlying meaning in the logic representation for the recordable code would be the same as shown in Figure 5b.

Hence, we have a unified representation in which the finding corresponds to a qualified observable and vice versa.^e The finding can be coded either as the code “blood

^dSNOMED CT January 2008 content contains exactly one right-identity statement and this is unrelated to procedures or anatomy. It states that the active ingredient of a direct substance is an active ingredient of the whole.

^eNote that the relation, hasJudgedLevel, indicates that the assigned qualifier is a qualitative clinical judgement as opposed to a simple quantitative value reading. The issue of representing and reasoning over absolute quantitative values and normative thresholds belongs to another paper.

Operation THAT site SOME (Lung OR is_part_of SOME Lung)
 Removal THAT site SOME Lung
 Removal THAT site SOME (Lobe THAT is_part_of SOME Lung)

Figure 3. Representation of operations on Lung and/or Parts of Lung.

pressure" (an observable) plus the qualifier "increased" or, if used frequently, as a single code for "increased blood pressure".

The notion of "finding" and "observable", properly understood, are "meta" to the ontology proper. "Findings" are things whose mere presence carries information; "observables" are things that must be qualified or given a value to convey information. There are large sections of SNOMED's ontology where the distinction between findings and observables follows the natural hierarchies. For example, disorders are generally findings and laboratory tests (or, rather, the physiological parameters underlying them) typically observables. However, there are cases where the distinction corresponds less well to natural clinical categories. For example, by these definitions, some physical signs are "findings"—e.g., the "presence of a lump"—others are "observables"—e.g., "pulse rate" or "body temperature."

Modularization

The national versions of SNOMED comprise an international core and national extension. Nevertheless, as an ontology it is currently classified and managed as a whole. Ontologies in OWL have an import mechanism by virtue of which it is easy to create ontologies comprising several modules. Although most tools present OWL entities as if they were unitary objects, in fact an OWL ontology consists simply of a set of axioms about those entities. Logic is "monotonic," additional axioms can lead to additional inferences, but they cannot annul previous inferences. (The inference that a class is inconsistent—"unsatisfiable"—is simply one more inference, although it generally indicates an error.) Hence, as well as adding new entities, OWL modules can add to the definitions of existing entities by adding new axioms about them. The result is a powerful method for composition and localization. Modern ontology editors, such as Protégé, support the user in the maintenance of modular ontology structures and automatically infer and

A

24184005|Finding of increased blood pressure (finding) →
 38936003|Abnormal blood pressure (finding) AND
 roleGroup SOME
 (363714003|Interprets (attribute) SOME
 75367002|Blood pressure (observable entity))

B

12763006|Finding of decreased blood pressure (finding) →
 392570002|Blood pressure finding (finding) AND
 roleGroup SOME
 (363714003|Interprets (attribute) SOME
 75367002|Blood pressure (observable entity))

Figure 4. A. SNOMED Representation for increased blood pressure. B. SNOMED Representation for decreased blood pressure.

A

Finding of increased blood pressure ↔
 Blood pressure THAT hasJudgedLevel SOME increased

B

Situation THAT includes

Blood pressure THAT hasJudgedLevel SOME increased

Figure 5. A. OWL Definition of Finding of increased blood pressure. B. Representation of the Recordable Concept for both the "Finding of increased blood pressure" and the observable "Blood pressure" qualified by "increased".

resolve "import graph" indirections when editing or reasoning over the ontology as a whole.

The reasons for breaking up the contents of an ontology into modules are numerous. Firstly, just as with chapters in books, or packages and modules in software, OWL modules allow the contents of an ontology to be partitioned into sub-units that are easier to maintain and exchange than would be a monolithic ontology.

Secondly, the modules mechanism of OWL affords a natural way to import content, particularly metadata, from other sources—e.g., specialized vocabularies for different realms or termsets for alternative languages.

Thirdly, modules can be used so that only the relevant portions of the ontology needed to be loaded for any given application. Statistical records published by The Health Improvement Network^f of the use of the READ codes in the UK show that only one thousand (or 1.1%) of all READ codes then available accounted for 81% of all coded information recorded by UK primary care practitioners between June 1st 2006 and May 31st 2007—and 10,000 codes (11%) accounted for 99% of it. These figures suggest that the vast majority of actual coding usage within a particular clinical subspecialty will be restricted to a tiny subset of the available coding content. There is little reason to expect the results to be radically different for SNOMED-CT, and preliminary evidence such as the well-publicized Kaiser Permanente/VA subset^g suggest similar results—that, for most applications, it should be possible to identify relatively small subsets of SNOMED that satisfy 95% or more of the requirements for a given use case or application. Such subsets could provide significant performance improvements if systems using SNOMED were first able to work primarily with a small subset but be able to load additional modules as necessary. This is particularly relevant for contexts in which post-coordination and, therefore, classification is required. (Note that techniques have been developed to support 'incremental classification' of additional modules that are loaded on an as needed basis.)¹⁴

Fourthly, national, local and specialized extensions of SNOMED could be formulated as add-on modules consisting primarily of definitions built from a core SNOMED vocabulary. Such so-called "pre-post-coordinated" extensions offer the advantage of limiting the number of concepts

^fTHIN: <http://www.thin-uk.com/>.

^ghttp://www.nlm.nih.gov/research/umls/Snomed/snomed_problem_list.html.

that have to be enumerated and named in advance while still tailoring the system to the needs of particular circumstances.

For example, it is neither practical nor necessary to enumerate named concepts for the family history of all possible diseases. In the modern world of “omics” and translational medicine, the family history of many more conditions is becoming relevant, so that the number of “pre-post-coordinated” concepts required is likely to grow rapidly. A special module of such named “pre-post-coordinated” family history concepts appropriate to a given genomics clinic or project could speed coding without affecting the logical content of the overall coding system.

Discussion

How Much OWL is Required?

The previous sections have suggested various extensions to the SNOMED schemas that are possible using OWL and would address identified weaknesses within SNOMED's current representation. Features of OWL required include conjunction, disjunction, full negation, existential restrictions, property chains, and general inclusion axioms. In addition, it is suggested that the ease with which OWL can be broken into modules would greatly simplify managing SNOMED's international and national extensions or subsets, and their subsequent localization. Note that we are not advocating the use of numerous other constructs in OWL. In particular, the proposed schemas limit the use of universal and maximum cardinality restrictions, which are known to reduce performance.

Schulz³ suggested using a fragment of OWL with attractive computational properties, EL++¹⁵ to achieve many of the reformulations suggested here. EL++ meets the above requirements with the exception of full negation and disjunction. However, we argue that these features are needed to address context and partonomy cleanly. While we recognize that a move to EL++ would allow important steps forward, we argue that the absence of negation and disjunction are serious disadvantages.

Is OWL Ready for SNOMED?

The main argument for using a less expressive representation such as EL++ is performance and scaling. The key issue is, therefore, whether it is plausible that a more radical reformulation using a larger subset of OWL would be practical. Do the reasoners and tools available for OWL 1.1 have the computational power and maturity to deal with an ontology of about 434,000 classes and more than a million relationships between them?

SNOMED can be loaded into Protégé^{4h}, the standard editor for OWL ontologies, and classified by FaCT++¹ without any extraordinary hardware requirements, provided a 64bit Java implementation is available. On a machine with a 2Ghz dual core CPU and 2 GB memory, classification takes 30 minutes; on a larger quad-core machine with 16GB memory something less than half this.

This situation is likely to improve rapidly due to new developments in three directions: faster reasoners, module-aware reasoners, and incremental reasoners. Hermit,¹⁶ a novel hypertextable-reasoner, supports the required expressivity and shows promise to deliver significantly more performance than FaCT++ today. IBM's SHER project¹ has produced a modularization strategy for the OWL-reasoner Pellet that allows it to classify all of SNOMED-CT effectively, and to classify several million individuals against the classified ontology in an acceptable time.¹⁷ In addition, incremental classifiers are being developed for both FaCT++ and Pellet, so that the need to reclassify the ontology from scratch can be minimized.

Note also that many uses of SNOMED for post coordination require only querying against a fixed classified ontology rather than reclassification, or even incremental classification. Such queries, which do not persist in the knowledge base, are much faster than classifications, which do persist.

Finally, as noted earlier, there exist relatively small subsets of SNOMED that are likely to satisfy the vast majority of uses cases.

In summary, we have established that the existing schemas can be classified using reasoners for more expressive description logics. Initial indications are that the more expressive schemas proposed in this paper will also scale using these same reasoners, but this remains to be tested. A feasibility study to answer this question should be a priority for the SNOMED community.

Is SNOMED Ready for OWL?

The need for significant quality assurance and development of SNOMED is widely accepted. The issues are increasingly well documented, and the complexity of using SNOMED in practice with HL7⁷ or Archetypes¹⁸ are increasingly well appreciated, although the methods to address them are not yet satisfactory. The main barrier to wholesale migration of SNOMED content to an OWL environment is likely to be the many errors and irregularities in the existing SNOMED content, a difficulty compounded by the sheer size of SNOMED. More manageable initial experiments in migration and curation of SNOMED CT content within native OWL representation and tooling may be possible if restricted to high value subsets.

Migration and Legacy

The pragmatics of migration of any large software artefact should not be under-estimated. In the case of SNOMED, it includes not only the ontology itself, but also the applications and metadata already based on, and in some cases erroneously embedded within, its current structure. One important but currently open question is how any migration of the SNOMED ontology to an OWL environment should treat that part of SNOMED's content that is derived from external classifications, such as ICD, for example: “172915004|Functional endoscopic sinus surgery—diagnostic endoscopy of nose or sinus NOS|”, or “371906007|Thrombolytic agent administered between 6 hours and 7 days before percutaneous coronary intervention (procedure)|” or “47198001|Cortex laceration with open intracranial wound AND prolonged loss of consciousness (more than 24 hours)

^hAvailable from protege.stanford.edu.

¹Personal communication, Dmitry Tarkov, 2008.

¹http://domino.research.ibm.com/comm/research_projects.nsf/pages/iaa.index.html.

AND return to pre-existing conscious level (disorder) |". The use cases and requirements need to be examined carefully to determine how much, if indeed any, of the content of such expressions should be modelled or whether they are rather best treated as instructions to coders and represented in some other way.

The issue of representing terms from external classifications aside, from the perspective of health-care institutions or third-party software developers, changes could be minimized. The basic set of codes, class names, and the delivery file structures for the pre-coordinated terms would remain largely unchanged. Hence, the changes might be comparable to a regular update, although perhaps larger in scope.

Outside the context dependent branch of the hierarchies, the changes advocated in the section on "Situations with explicit context" can be accomplished by 'wrapping' the relevant definitions into the expression "Situation that includes . . .", an operation that can be automated easily. The changes to the handling of partonomy would have little effect on applications and end users, although they would greatly simplify the work of developers.

The changes we propose to the context dependent branch are significant: many concepts would need to be redefined to the extent that a new code was felt to be required, and the guidelines for using context dependent codes would have to be revised. However, despite recent less radical efforts to improve SNOMED's context dependent codes, they are known to remain problematic. Further, they are modest in number. A thorough revision that produced a uniform structure integrated with the rest of the SNOMED structure rather than a piecemeal revision with multiple variants should be attractive.

The exact costs of implementing all the recommendations expressed in the present paper cannot be determined with precision because of the irregularities of SNOMED. Where the structure of SNOMED is sound and regular, it could be done largely via scripting. Where the structure is flawed, formal classification by the OWL reasoner of a scripted transform would make those flaws more obvious, but manual detection and revision would still be required. Where the naming is regular, this work can be aided by lexical techniques.¹⁹ However, other results^{18,20} suggest that lexical methods are far from reliable as do experiments with inter-rater variability amongst SNOMED coders.¹⁹ Furthermore, having to depend on lexical techniques rather than verifiable models and logical inference to support software interoperability and critical clinical systems calls into question SNOMED's claim to be a "reference terminology."

Summary

In this paper we argue the case for reformulating the definitions of classes in SNOMED's stated form so as to utilize constructors available in the ontology standard OWL 1.1, a considerable extension of SNOMED's current underlying formalism Ontolog. The main advantages of this reformulation are: (1) a simpler, uniform representation of situations with explicit context, (2) a more flexible way of handling definitions of classes, (3) a simple and uniform way to specify partonomic information for procedures and findings, (4) a clearer and more principled relation between

observable entities and findings, and (5) the chance to modularize SNOMED. The overall result would, almost certainly, be a representation that scaled better cognitively—i.e., that was easier to use, easier to maintain, simpler to query, and easier to use in the development of associated software.

The relationship between codes and EHR models such as the HL7 RIM, Archetypes, and CEN 13606 has been discussed in a previous paper.⁴ A cleaner formulation as advocated here would likewise facilitate better specification of the binding between terminology and information models for messages and EHRs. These patterns also integrate much more easily with the OBO family of ontologies and other ontologies being used in molecular biology than do SNOMED's current schemas.

The price to pay for this reformulation is three-fold. Firstly, the syntax would need to be transformed to OWL 1.1; this is an entirely automatic process already catered for by existing tools. More significantly, to take full advantage of the proposed schemas, the explicit content of SNOMED's definitions would need to be reviewed and extended. However, that could be done progressively. Thirdly, different tools and classifiers would need to be used. This would require some re-engineering of tools to cope with the scale of SNOMED, and the classification time would inevitably increase somewhat. Current indications, however, are that the effort on tools and the increase in classification time could both be kept modest. However, but this remains to be proven.

Indeed, all of these assertions need to be tested by a feasibility study on a limited subset of SNOMED. A number of modest sized—on the order of 25,000 concept—subsets of SNOMED exist that could be used for feasibility tests. The transfer of responsibility to an international standards body (the IHTSDO) provides a natural point to consider new developments. Although increasing, use of SNOMED is still in its early stages, even in the United Kingdom. Changes made now will be much easier to implement than changes made in the future when the legacy is greater. The result of the suggested changes would be a simpler and clearer representation. Why do it the hard way? At a minimum, the feasibility of alternative schemas should be tested.

References ■

1. Bodenreider O, Smith B, Kumar A, Burgun A. Investigating subsumption in SNOMED CT: An exploration into large description logic-based biomedical terminologies. *Artif Intell Med* 2007;39:183–95.
2. Schulz S, Hanser S, Hahn U, Rogers J. The semantics of procedures and diseases in SNOMED CT. *Methods Inform Med* 2006;45:354–8.
3. Schulz S, Suntisrivaraporn B, Baader F. SNOMED CT's Problem List: Ontologists' and logicists' therapy suggestions. 2007; *Medinfo* 2007: IOS Press; 802–806.
4. Rector AL. What's in a code: Towards a formal account of the relation of ontologies and coding systems. 2007; *Medinfo* 2007: Brisbane, Australia: IOS Press; 730–4.
5. Rector A, Qamar R, Marley T. Binding ontologies & coding systems to electronic health records and messages. 2006; *Formal Biomedical Knowledge Representation (KR-MED 2006)* (To ap-

- pear in J Applied Ontologies) CEUR Workshop Proceedings 222: Baltimore: CEUR; 11–19.
6. Horridge M, Drummond N, Goodwin J, Rector A, Stevens R, Wang H. The Manchester OWL syntax. 2006; OWL: Experiences and Directions (OWLED 06): Athens, Georgia: CEUR Available at: http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-216/submission_9.pdf. Accessed September 2008.
 7. Cheetham E, Markwell D, Dolin R. Using SNOMD CT in HL7 Version 3: Impemenation Guide, Release 1.3. HL7 Ballot Document. 2006.
 8. Rector AL, Bechhofer S, Goble CA, Horrocks I, Nowlan WA, Solomon WD. The GRAIL concept modelling language for medical terminology. *Artif Intell Med* 1997;9:139–71.
 9. Patrick J. Metonymic and holonymic roles and emergent properties in the SNOMED CT Ontology. 2006; The Australian Ontology Workshop (AOW 2006): Hobart, Tasmania, Australia.
 10. Johansson I. On the transitivity of parthood relations. In: H. Hochberg, K. Mulligan (eds). *Relations and Predicates*. Frankfurt: Ontos Verlag; 2004, pp 161–81.
 11. Rector A, Rogers J, Bittner T. Granularity, scale & collectivity: When size does and does not matter. *J Biomed Inform* 2006;39: 333–49.
 12. Schulz S, Hahn U. Towards a computational paradigm for biomedical structure. <http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-102/>; 2004; KR 2004 Workshop on Formal Biomedical Knowledge Representation: Whistler, Canada: CEUR; 63–71.
 13. Schulz S, Hahn U. Mereotopological reasoning about parts and (w)holes in bio-ontologies. 2001; *Formal Ontology in Information Systems (FOIS-2001)*: Ogunquit, ME: ACM; 210–21.
 14. Grau BC, Halaschek-Wiener C, Kazakov Y. History matters: Incremental ontology reasoning using modules. 2007; *Proc 6th International Semantic Web Conference (ISWC-2007)*: Busan, Korea: Springer, LNCS 4825.
 15. Baader F, Lutz C, Suntisrivaraporn B; CEL—A polynomial time reasoner for life science ontologies. 2006; *3rd International Joint Conference on Automated Reasoning (IJCAR'06)*: Springer-Verlag (LNI 4130);287–91.
 16. Motik B, Shearer R, Horrocks I. Optimized reasoning in description logics using hypertableaux. 2007; *21st Conference on Automated Deduction (CADE-21)*: Bremen, Germany: Springer LNAI; 67–83.
 17. Patel C, Cimino J, Dolby J, et al.; Matching patient records to clinical trials using ontologies. 2007; *International Semantic Web Conference (ISWC 2007)*.
 18. Qamar R, Rector A. Unambiguous data modeling to ensure higher accuracy term binding to clinical terminologies. 2007; *Medinfo 2007*: Brisbane, Australia: 675–8.
 19. Chiang MF, Hwang JC, Yu AC, Casper DS, Cimino JJ. Reliability of SNOMED-CT Coding by Three Physicians using Two Terminology Browsers. 2006; *AMIA Annual Symposium*: 131–5.
 20. Qamar R. Semantic Mapping of Clinical Model Data to Biomedical Terminologies to Facilitate Interoperability [dissertation]. Manchester UK: University of Manchester; 2007.